arXiv:1612.06524v1 [cs.CV] 20 Dec 2016

3D Human Pose Estimation = 2D Pose Estimation + Matching

Ching-Hang Chen
Carnegie Mellon University
chinghac@andrew.cmu.edu

Deva Ramanan
Carnegie Mellon University
deva@cs.cmu.edu

Abstract

We explore 3D human pose estimation from a single RGB image. While many approaches try to directly predict 3D pose from image measurements, we explore a simple architecture that reasons through intermediate 2D pose predictions. Our approach is based on two key observations (1) Deep neural nets have revolutionized 2D pose estimation, producing accurate 2D predictions even for poses with self-occlusions (2) "Big-data"sets of 3D mocap data are now readily available, making it tempting to "lift" predicted 2D poses to 3D through simple memorization (e.g., nearest neighbors). The resulting architecture is trivial to implement with off-the-shelf 2D pose estimation systems and 3D mocap libraries. Importantly, we demonstrate that such methods outperform almost all state-of-the-art 3D pose estimation systems, most of which directly try to regress 3D pose from 2D measurements.

1. Introduction

Inferring 3D human pose from image measurements is classic task in computer vision, dating back to the iconic work of Hogg [9] and O'Rourke and Badler [19]. Such a technology has immediate applications in various tasks such as action understanding, surveillance, human-robot interaction, and motion capture, to name a few. As such, it has a long and storied history. We refer the reader to various surveys for a broad overview of the popular topic [6, 16].

Previous approaches often make use of a highly sensored environment, including video streams [33, 26], multiview cameras [3, 8], depth images [20, 32, 24]. In this work, we focus on the "pure" and challenging setting of recovering 3D body pose with a single 2D RGB image [15, 31, 22, 28].

Our key insight to the problem is leveraging recent advances in 2D image understanding, made possible through the undeniable impact of deep learning. While originally explored for coarse recognition tasks such as image classification, recent methods have extended such network architectures to "fine-grained" human pose estimation, where the task is formulated as one of *2D heatmap prediction* problem [29, 18, 27, 10]. One of the long standing challenges in 2D human pose estimation has been estimating poses under self-occlusions. Difficulties with occlusions has a recurring motivation for reasoning in 3D rather than 2D. But one of our conclusions is that state-of-the-art 2D methods do a surprisingly good job of pose estimation under occlusion. Given this observation, the remaining challenge remains lifting 2D poses to 3D.

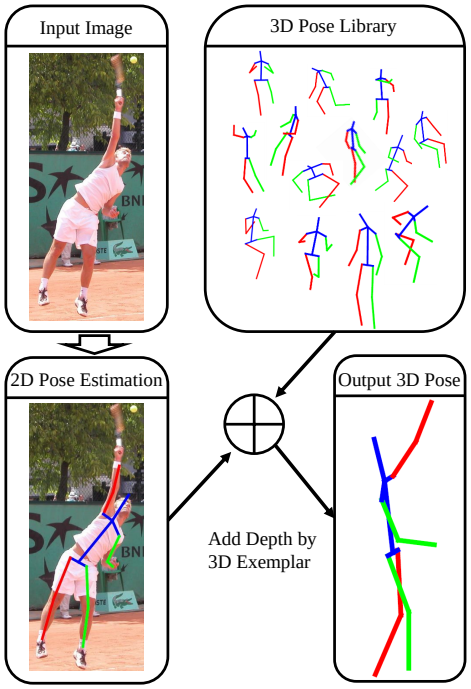

Figure 1. The overview of proposed framework. Given an input image, first perform 2D pose estimation, then perform feature matching with 3D pose library to extract exemplar. Finally 3D pose is completed with depth taken from exemplar.

Inferring 3D structure from 2D correspondences has also been a storied problem in computer vision, often addressed in multiview setting as structure from motion. In the context of monocular human pose estimation, the relevant cues seem to be semantic rather than geometric. One can es-

timate 3D postures from a 2D skeleton based on high-level knowledge derived from anthropometric, kinematic, and dynamic constraints. Inspired by the success of data-driven architectures, we explore a simple non-parametric encoding of such high-level constraints: given a 3D pose library, we generate a large number of 2D projections (from virtual camera views). Given this training set of paired (2D,3D) data and predictions from a 2D pose estimation algorithm, we report back the 3D pose associated with the closet matching 2D example from our library. Our entire pipeline is summarized in Fig. 1. The available large-scale 3D human-pose datasets are collected by the MoCap system, which is installed under a constrained environment. Due to the low diversity image content (in terms of background and subject), the end-to-end learning scheme is not preferable. Our approach tackles 3D postures independently from input images, thus, could apply on images in the wild!

Evaluation: We perform an extensive evaluation of our method on widely benchmarked 3D human-pose datasets, including Human3.6M [11]). We follow standard train/test protocol splits. but our analysis reveals that there has been inconsistent reporting in the literature, both in terms of test sets and evaluation criteria. To make our results as transparant as possible, we report performance for all metrics and splits we could find. One of our surprising findings is the impressive performance of our simple pipeline: we outperform essentially all prior work on all metrics. Our entire pipeline, even given the non-parametric matching step, returns a 3D pose given a 2D image in under 1 second (160ms for 2D estimation by CNN, 26ms for exemplar matching under 200,000 data size). Finally, to promote future progress, we perform an exhaustive analysis of additional baselines with upper bounds that reveal the continued benefit of working with intermediate 2D representations and data-driven encoding of 3D constraints.

2. Related Work

Here we review related works on 3D human pose prediction most relevant to our approach.

(Deep) Regression: Most existing work that makes use of deep features tends to formulate the problem as a direct 2D image to 3D pose regression task. Li *et al.* [15] use deep learning to train a regression model to predict 3D pose directly from images. Tekin *et al.* integrate spatio-temporal features via an image sequence to learn regression model for 3D pose mapping. We conduct experiment showing that it is worth exploring 2D pose as a core intermeedate representation.

Intermediate 2D Pose: Other approaches have explored pipelines that use 2D poses as an intermediate result. Most focus on the second-stage that lifts 2D estimates to 3D. This is classically treated as a constrained optimization problem who's objective minimizes the 2D reprojection error of an unknown 3D pose and unknown camera [33, 28, 21, 2]. The optimization problem is often subject to kinematic constraints [30, 25], and sometimes 3D poses are assumed to live a in low-dimensional subspace to better condition the optimization. Such optimization-based approaches could be senstive to initialization and local minima, and often require expensive constrained solvers. We use a simple data-driven matching algorithm that is quite fast. Importantly, our brute-force matching returns a score for each putative match, which can be normalized into a *distribution* of 3D poses.

Exemplar based: Previous work has also explored example-based methods, dating back at least to work of [23]. One of the central challenges is generalization to novel poses outside the training set. [12] propose matching upper and lower bodies individually, to allow for novel compositions at test-time. [31] adapt examplars to better match image measurements with an energy minimization approach. [22] synthesize new 2D images with image-based rendering. Other methods also warp 3D exemplars to 2D image descriptors, classically shape-context [1, 17] or silhouette features [5]. In our work, we show that a modest number of exemplars (200,000), combined with a trivial "closed-form" algorithm for warping a 3D exemplar to exactly project to 2D pose estimates, outperforms more complex methods.

3. Approach

In this section, we describe our method for estimating 3D human pose given a single RGB image. We make use of a probabilistic formulation over variables including the image I, the 3D pose $\mathbf{X} \in \mathbb{R}^{N\times 3}$, and the 2D pose $\mathbf{x} \in \mathbb{R}^{N\times 2}$, where N is the number of articulated joints. We write the joint probabililty as:

$$p(\mathbf{X}, \mathbf{x}, I) = p(\mathbf{X}|\mathbf{x}, I) \cdot p(\mathbf{x}|I) \cdot p(I) \tag{1}$$

where the above makes no limiting assumptions by itself.

Conditional independence: Let us now assume that the 3D pose $\mathbf{X}$ is *conditionally independent* of image I given the 2D pose $\mathbf{x}$. This is equivalent to the implication that given a 2D skeleton, the prediction of its corresponding 3D skeleton would not be affected by 2D image measurements. While this is not quite true (the appearance of a frontal versus profile face likely has implications for 3D), it seems to be a reasonable first-order approximation. Moreoever, this factorization still allows for $p(\mathbf{x}|I)$ to be arbitrarily complex (likely needed to accurately model complex dependencies between 2D projections and image features during occlusions). Given this conditional independance, one can write:

$$p(\mathbf{X}, \mathbf{x}, I) = \underbrace{p(\mathbf{X}|\mathbf{x})}_{\text{NN}} \cdot \underbrace{p(\mathbf{x}|I)}_{\text{CNN}} \cdot p(I) \tag{2}$$

We tackle the second term with a image-based CNN that predicts 2D keypoint heatmaps. We tackle the first term with a non-parametric nearest-neighbor (NN) model. We describe each term in turn below.

3.1. Image-Based 2D Pose Estimation

Given the above independance assumption, we would first like to predict 2D pose given image measurements. We model the conditional of 2D pose given an image as

$$P(\mathbf{x}|I) = CNN(I) \tag{3}$$

where we assume CNN is a nonlinear function that returns N 2D heatmaps (or marginal distributions over the location of individual joints). We make use of convolutional pose machines (CPMs) [29], which return precisely N heatmaps for individual body joints. We normalize the heatmaps so that they can be interpreted as marginal distributions for each joint. CPM is a near-state-of-the-art pose estimation system (88.5% PCKh on MPII dataset [4], quite close to the state-of-the-art value of 90.9% [18]). Note the off-the-shelf CPM model was trained on MPII dataset, which is a somewhat limited dataset in that annotations are provided through manual inspection. We fine-tune this model on the large scale Human3.6M [11] training set, which contains annotations acquired by a Mocap system (allowing for larger-scale labeling).

3.2. Nonparametric 3D shape model

We model $P(\mathbf{X}|\mathbf{x})$ with a non-parametric nearest neighbor model. We will follow a notational convention where $\mathbf{X} = [X, Y, Z]$ and $\mathbf{x} = [x, y]$. Assume that we have library of 3D poses $\{\mathbf{X}_i\}$ paired with a particular camera projection matrix $\{M_i\}$, such that the associated 2D poses are given by $\{M_i(\mathbf{X}_i)\}$. If we want to consider multiple cameras for a single 3D pose, we add another copy of the 3D pose with a different camera matrix to our library. We define a distribution over 3D poses based on reprojection error:

$$P(\mathbf{X} = i|\mathbf{x}) \propto e^{-\frac{1}{\sigma^2}||M_i(\mathbf{X}_i)-\mathbf{x}||^2} \tag{4}$$

where the MAP estimate is given by the 1-nearest neighbor (1NN). We explore two extensions to the above basic framework.

Virtual cameras: We can further reduce the squared reprojection error by searching over small pertubations of each camera. This involves solving a camera resectioning problem [7], where an iterative solver can be initialized with M_i:

$$M_i^* = \underset{M}{\operatorname{argmin}} ||M(\mathbf{X}_i) - \mathbf{x}||^2 \tag{5}$$

In practice, we construct a shortlist of k candidates that score well according to (4), and resort them according to optimal camera matrix. We found that optimizing over cameras produced a small but noticeable improvement in our experiments. Unless otherwise specified, we choose $k = 10$ in our experiments.

Warped exemplars: Much previous work on exemplars introduce methods for warping exemplars to better match the 2D pose estimates, often formulated as an inverse kinematics optimization problem. We describe an extremely lightweight method for doing so here. We first align the 3D exemplar to the camera-coordinate system used to compute the projection $\mathbf{x}$. This is done with a 3D rigid transformation given by the camera extrinsics encoded in M_i (or M_i^*). In practice we use a training set $\{\mathbf{X}_i\}$ where 3D exemplars are already aligned to their projections $\{\mathbf{x}_i\}$, implying that extrinsics in M_i reduce to an identity matrix (which is the case for the Human3.6M dataset [11], since 3D poses are specified in camera coordinates of their associated image projections). Given this alignment, we simply replace the (X_i, Y_i) exemplar coordinates with their scaled 2D counterparts (x, y) under a weak perspective camera model:

$$\mathbf{X}_i^* = \begin{bmatrix} sx & sy & Z_i \end{bmatrix}, \quad \text{where} \quad s = \frac{\text{average}(Z_i)}{f} \tag{6}$$

where f is the focal length of the camera (given by the intrinsics in M_i) and average(Z_i) is the average depth of the 3D joints. We demonstrate that this simple warping method rivals much more complex energy minimization methods (Fig. 4).

4. Experiments

In our experiments, we test a variety of variations of our proposed pipeline.

Qualitative results: We first present some qualitative results. Fig. 2 shows results on challenging examples from subject S11 of Human3.6M. We choose examples with self-occlusions and sitting poses. To demonstrate the accuracy of the 3D predictions, we visualize novel viewpoints. We then apply the proposed method on Leeds Sports Pose(LSP) dataset [13] to test cross-dataset generalization. We posit that our pipeline will generalize across image variation (due to the underyling robustness of our 2D pose estimation system) but maybe limited in the 3D estimates due to the library used (from Human3.6M). Fig. 8 shows that our warping approach produces plausible 3D poses even when the activity class is not included in Human3.6M. This implies that our method can reliable estimate 3D poses in the wild!

4.1. Quantiative Evaluation Protocols

We use Human3.6M for quantitative evaluation and analysis. It appears that multiple train/test splits have been used in the literature, as well as different interpretations of mean per joint position error (MPJPE). We summarize them here.

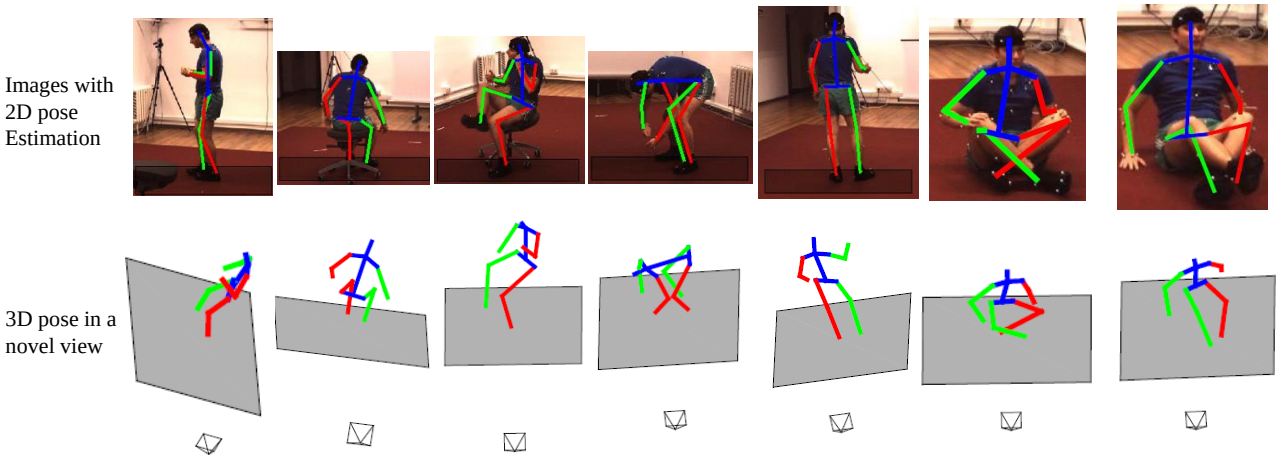

Figure 2. Qualitative results of subject S11 from Human3.6M, examples with self occlusion and challenging pose are selected.

Protocol 1: In works [31, 14, 22], the entire dataset were partitioned into six training subjects (S1, S5, S6, S7, S8, S9), and one testing subject (S11). Evaluation is performed on every 64th frame of S11's video clips. In this configuration, there are total 1.8 million 3D poses available in the training set. MPJPE between the ground truth 3D pose and the estimated 3D pose by first aligning poses with a rigid transformation [14].

Protocol 2: Others [33, 26, 15] use five subjects (S1, S5, S6, S7, S8) for training, and two subjects (S9, S11) for testing. We follow [33]'s setup that downsamples the videos from 50 fps to 10 fps. Here, MPJPE is evaluated *without* a rigid transformation, following the original *h36m* protocol: both the ground-truth and predicted 3D pose is centered with respect to a root joint (*ie.* pelvis). In contrast to Protocol 1, this evaluation can be sensitive to a single poorly-predicted joint, particularly if it is the root [11].

To compare to published performance numbers, we use the appropriate protocol as needed. From our own experience, we find Protocol 1 to be simpler and more intuitive, and so focus on it for our diagnostic evaluations.

4.2. Comparison to state-of-the-art (Protocol 1)

Final system: Table 1 compare MPJPE for each activity class. Our approach clearly outperforms [31] and [22]. ("Ours" in the tables of comparison throughout the experiment refers to the warped exemplar $\mathbf{X}^*$ described in Section 3.2)

Performance given ground-truth 2D: A common diagnostic is evaluating performance given ground-truth 2D poses, written as *gt*. Table 2 shows that our simple matching + warping outperforms [31], who use a complex iterative appearance for matching and warping exemplars to image evidence. Our diagnostics will later show that even matching exemplars *without* warping outperforms prior art, indicating the remarkable power of a simple NN baseline.

Method	2D source	3D source	MPJPE
Yasin [31]	64,000k	380k	108.3
Rogez [22]	207k	190k	88.1
Ours	18k	18k	85.94
Ours	180k	180k	**82.37**
Ours	1,800k	1,800k	82.72

Table 3. Comparison under different training sizes with [31] and [22]. Our approach yields the best performance at the source size of 180k.

Size of trainset: Table 3 shows the MPJPE versus the training data size. Since approaches deal with 2D source and 3D sources differently, we list both sizes. Yasin *et al.* [31] project multiple 2D poses from each 3D exemplar (with virtual cameras) to create abundant 2D poses for matching, and Rogez *et al.* [22] synthesized 2D images to boost performance of learned classifier. Our approach makes use of the default training data, where each 3D pose is paired with a single 2D projection. As can be seen, our approach reaches the best performance using a pose library of 180k 3D-2D pairs, and results in a lower MPJPE even given 18k. The slight increase in MPJPE for large training sets seems to be related to noise from 2D pose estimation, since we observe a monotonically decrease when ground truth 2D poses are given (Fig. 6).

4.3. Comparison to state-of-the-art (Protocol 2)

Final system: Table 4 provides the comparison to [33] and [26] using Protocol 2. Note that in both these works, temporal smoothness was exploited by taking a short image sequences as input. Even though we do not use temporal information, our system is quite close to state-of-the-art. A qualitative comparison to [33] is also provided in Fig. 3.

Method	Direction	Discuss	Eat	Greet	Phone	Pose	Purchase	Sit	SitDown
Yasin [31]	88.4	72.5	108.5	110.2	97.1	81.6	107.2	119.0	**170.8**
Rogez [22]	-	-	-	-	-	-	-	-	-
Ours	**71.63**	**66.60**	**74.74**	**79.09**	**70.05**	**67.56**	**89.30**	**90.74**	195.62
Method	Smoke	Photo	Wait	Walk	WalkDog	WalkPair	Avg	Median	-
Yasin [31]	108.2	142.5	86.9	92.1	165.7	102.0	108.3	-	-
Rogez [22]	-	-	-	-	-	-	88.1	-	-
Ours	**83.46**	**93.26**	**71.15**	**55.74**	**85.86**	**62.51**	**82.72**	**69.05**	-

Table 1. Comparison to [31] by **Protocol 1**. Our results are clearly state-of-the-art. Please see text for more details.

Method	Direction	Discuss	Eat	Greet	Phone	Pose	Purchase	Sit	SitDown
Yasin [31]	60.0	54.7	71.6	67.5	63.8	61.9	55.7	73.9	110.8
$\mathbf{X}^*\|gt$ (Ours)	**53.27**	**46.75**	**58.63**	**61.21**	**55.98**	**58.13**	**48.85**	**55.60**	**73.41**
Method	Smoke	Photo	Wait	Walk	WalkDog	WalkPair	Avg	Median	-
Yasin [31]	78.9	96.9	67.9	47.5	89.3	53.4	70.5	-	-
Ours	**60.25**	**76.05**	**62.19**	**35.76**	**61.93**	**51.08**	**57.50**	**51.93**	-

Table 2. Comparison to [31] by **Protocol 1** given 2D ground truth.

Method	Direction	Discuss	Eat	Greet	Phone	Pose	Purchase	Sit	SitDown
Zhou [33]	**87.36**	109.31	**87.05**	**103.16**	116.18	106.88	**99.78**	**124.52**	**199.23**
Tekin [26]	102.41	147.72	88.83	125.38	118.02	112.38	129.17	138.89	224.9
Ours	89.87	**97.57**	89.98	107.87	**107.31**	**93.56**	136.09	133.14	240.12
Method	Smoke	Photo	Wait	Walk	WalkDog	WalkPair	Avg	Median	-
Zhou [33]	107.42	139.46	118.09	79.39	114.23	97.70	**113.01**	-	-
Tekin [26]	118.42	182.73	138.75	**55.07**	126.29	**65.76**	124.97	-	-
Ours	**106.65**	**139.17**	**106.21**	87.03	**114.05**	90.55	114.18	93.05	-

Table 4. Comparison to [33] and [26] by **Protocol 2**.

Method	MPJPE
Zhou [33]$\|gt$, single frame	50.04
Zhou [33]$\|gt$, temporal smoothness	49.64
Zhou [33]$\|gt_{syn}$, single frame	48.08
Zhou [33]$\|gt_{syn}$, temporal smoothness	**47.57**
$\mathbf{X}^*\|gt$, $k = 1$	51.06
$\mathbf{X}^*\|gt$, $k = 10$	49.55

Table 5. Comparison to [33] by virtual camera fine-tuning (5) given ground truth 2D pose. k is the number of candidate exemplar extracted in the shortlist. Choose k equals 10 in our approach gives a lower error in the cases no 3D prior knowledge is accessible. When 3D prior is provided by synthesized input, it yields a better result.

Note that in the second example (the second row in Fig. 3), both approaches predict the wrong left and right arms.

Performance given ground-truth 2D: Our strong performance in Fig. 3 might be attributed to better 2D pose estimation. Therefore, we investigate the case given ground truth 2D pose, following Zhou's diagnostic protocol [33]: evaluate MPJPE up to a 3D rigid body transformation including scale, only on the first 30 seconds of the first camera in Human3.6M. For a fair comparison, we make use the same training set of 3D-2D training data for both methods. In [33], they also examined the case where 3D structure prior is accessible. It is done by providing 2D pose rendered in a novel camera view, denoted as gt_{syn}. It has been proved that 3D prior provides critical message in solving the spatial structure, however, is not obtainable in real cases. The results are shown in Table 5. With a shortlist of $k = 10$ matches, camera resectioning (5) and exemplar warping (6) produces a slightly lower error than [33]'s approach without 3D prior.

Finally, a qualitative comparison is provided in Fig. 4. Our approach produces much lower 2D reprojection error, while Zhou's method suffers from the restriction of 3D poses to a low-dimensional subspace.

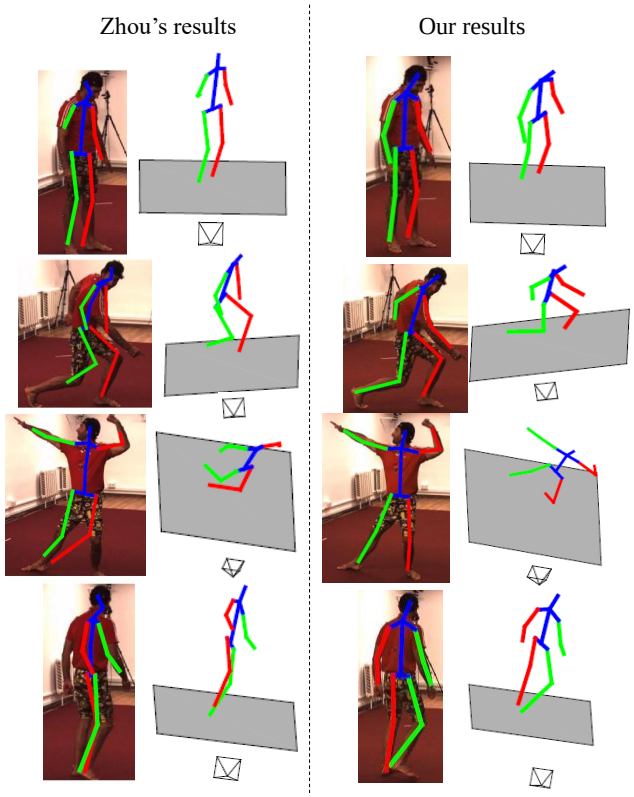

Figure 3. Qualitative comparison with the results of Zhou [33] in "Pose" class. The first two columns are input image with predicted 2D pose and 3D pose in a novel view of [33], the right two columns are our results of warped exemplar. Note the left and right arms are incorrectly estimated for both approaches in the second row example.

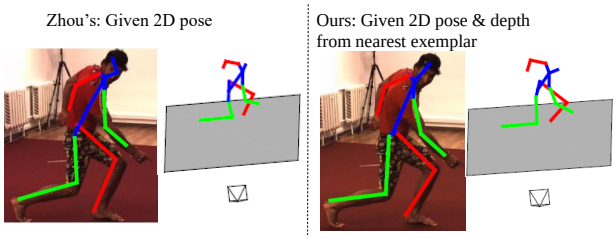

Figure 4. Qualitative comparison given ground truth 2D pose. The left side is Zhou's [33] result, the right side is our result with given 2D pose and the depth from nearest exemplar. Since Zhou combines 3D pose linearly from a set of basis, the 2D reprojection does not exactly equal to the given 2D, while we take given 2D pose directly. Also, in both cases, the estimated 3D poses seem to be plausible.

4.4. Diagnostics

We now perform an extensive set of diagnostics to reveal the strength of our individual components, as well as upper-bound analysis that is useful for guiding future work.

	Error (mm)	
Prediction	Avg	Median
$\mathbf{X}\|\mathbf{x}$	85.52	75.04
$\mathbf{X}^*\|\mathbf{x}$	**82.72**	**69.05**
$[s\mathbf{x}\ Z_{GT}]$	43.86	30.19

Table 6. Given the predicted 2D pose $\mathbf{x}$, warped exemplars $\mathbf{X}^*$ outperform unwarped exemplars $\mathbf{X}$ by a reasonable margin. An upper-bound for warped exemplars (that use (x, y) estimates from the predicted pose and z estimates from the ground-truth 3D pose suggests that significant further improvement is possible, even given existing 2D pose estimation systems.

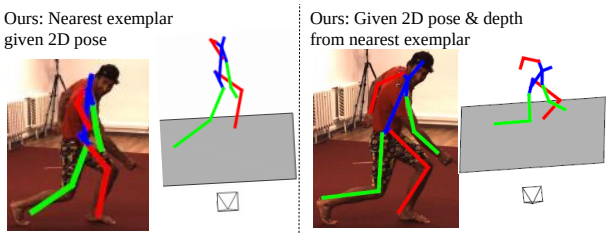

Figure 5. Qualitative comparison of warped and unwarped exemplars given ground truth 2D pose. The left side is exemplar $\mathbf{X}$ while the right side is warped exemplar $\mathbf{X}^*$. Clearly warping the exemplar to better match the 2D projection results in a more realistic pose.

Effect of warping: We evaluate the benefits of warping ($\mathbf{X}^*$ vs $\mathbf{X}$) in Table 6. The evaluation metric is based on **Protocol 1**. MPJPE is reported in both a average and median statistic, which is measured in mm. It is clear that warping exemplars $\mathbf{X}^*$ is a simple and effective approach to reducing error. Perhaps shockingly, even without warping, simply matching to a set of fixed 3D exemplar projections outperforms the state-of-the-art (see Table 1 & Table 6)! To analyze an upper-bound for our warping approach, we combine 2D estimates (x, y) with depth values Z given by the ground-truth 3D pose Z_{GT}. In the last row, the performance of combining ground truth depth Z_{GT} with $\mathbf{x}$ is listed as a reference baseline. This suggests that one can still dramatically lower error by 2X even when continuing to use the output of current 2D pose estimation systems.

Warping given ground-truth 2D: Next, we compute the error for the case that ground truth 2D pose is given, as shown in Table 7. We write $|gt$ to emphasize that methods now have access to 2D ground-truth pose estimates. We first note that matching *unwarped* examples rivals the accuracy of state-of-the-art (see Table 2 & Table 7). This again suggests the remarkable power a simple NN baseline based on matching 2D projections. That said, warping still improves results by a considerable margin. A qualitative example is provided in Fig. 5.

Prediction	Avg	Median
$\mathbf{X}\|gt$	70.93	65.35
$\mathbf{X}^*\|gt$	**57.50**	**51.93**

Table 7. Given ground truth 2D pose, the comparison of exemplar $\mathbf{X}$ and warped exemplar $\mathbf{X}^*$.

Prediction	Avg	Median
$\mathbf{X}\|GT$	60.11	55.36
$\mathbf{X}^*\|GT$	**37.32**	**33.91**

Table 8. We analyze performance given the optimal matching 3D training exemplar "GT" (in terms of 3D error wrt the ground-truth test 3D pose). Simply reporting this optimal match produces an error of 60mm, around 10mm lower than the actual match found given an ideal 2D pose-estimation system (Table 7). Warping this exemplar $X^*|GT$ significantly improves accuracy.

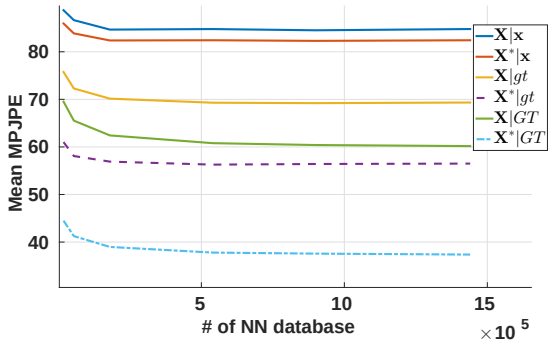

Figure 6. Mean MPJPE by **Protocol 1** versus database size. Each curve represents one configuration. In general, when the database size increases, the MPJPE decreases. Note the error stops decreasing after some point, at (2×10^5) for cases of 2D pose by CPM. In the case ground truth 2D or 3D is given, it converges at higher database size (5×10^5) when generalization is applied. The results explain that with more accurate 2D pose estimation, we can benefit more from large datasize.

Warping given optimal exemplar match: In addition, we are curious about the optimal performance possible given our training set of (3D,2D) pairs. We first compute the optimal exemplar that minimizes 3D reprojection error (up to a rigid body transformation) to the true 3D test pose. We write the index of this best match from the training set as $i = GT$. We would like to see the effect of warping given this optimal match. We analyze this combination in Table 8. This suggests that, in principle, error can still be significantly reduced (by almost 2X) even given our fixed library of 3D poses. However, it is not clear that this is obtainable given our pipeline because it may require image evidence to select this optimal 3D exemplar (violating the conditional independence assumption from (2)).

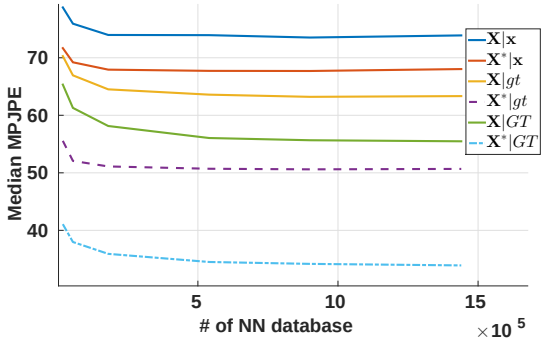

Figure 7. Median MPJPE by **Protocol 1** versus database size. The median error is commonly lower than mean error in Fig. 6, due to the evaluation metric in L2 loss, which has large penalty when some joints are incorrectly predicted (severely deviated from ground truth).

Effect of trainset size: An important aspect to investigate is the influence of database size. Here we investigate the error versus the number of exemplars in the database. We evaluate performance versus a random fraction of our overall database, using the experimental protocol in Section 4.4. Average MPJPE versus exemplar number is plotted in Fig. 6. As expected, more data results in lower error, though diminishing results are observed (even in log scale). This is reasonable since training data is extracted from videos captured at 50 fps, implying that correlations over frames might limit the benefit of additional frames. We see that convergence is also effected by the quality of the 2D pose estimates: error given ground-truth 2D poses plateaus at 5×10^5, while 2D pose estimates plateau even sooner at 2×10^5. We posit that a more restricted 3D pose prior (implicitly enforced by a small randomly-sampled 3D library) helps given inaccurate 2D pose estimates. But in either case, *exemplar-based 3D matching is effective even for modestly-size training sets (200000)*. This analysis appears to suggest that better 2D pose estimates are needed to take advantage of "bigger" 3D datasets.

Since the joint prediction error is not a normal distribution, we also plot median error in Fig. 7. We see that the median is generally lower than mean error, and the difference between the two becomes smaller when ground truth 2D or 3D is given. This may suggests that errors are often due to a single incorrect joint prediction, which would significantly impact average error but not the median.

5. Conclusion

We present an simple approach to 3D human pose estimation by performing 2D pose estimation, followed by 3D exemplar matching. The simplicity and efficiency of our method, combined with its state-of-the-art performance on

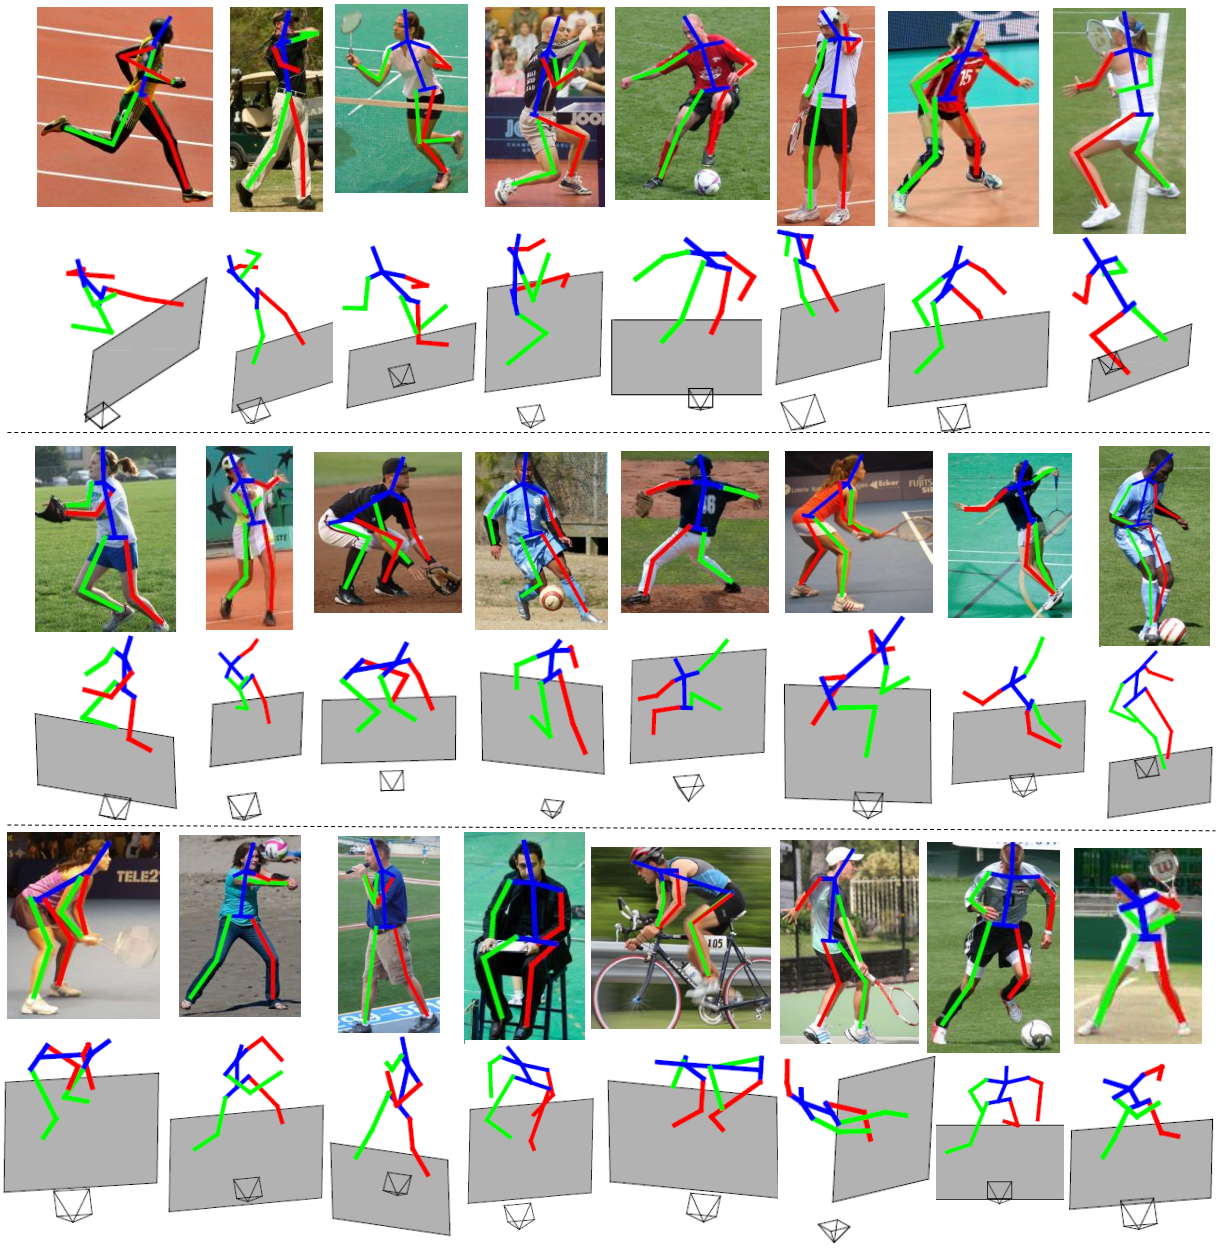

Figure 8. Qualitative results on LSP dataset [13].

both benchmark datasets and unconstrainted "in-the-wild" imagery, suggests that such simple baselines should be used for future benchmarking in 3D pose estimation. Our results also suggest that 3D inference is, in some sense, "all about 2D", at least in the case of articulated objects. Indeed, one of the surprising results of our analysis was the high performance of 2D pose estimation systems even under occlusions, suggesting that 2D estimates can infact be reliably estimated without directly reasoning about depth. Moreover, we find that 3D memorization, followed by simple 2D warping, is as effective strategy for imputing depth from 2D.

References

[1] A. Agarwal and B. Triggs. 3d human pose from silhouettes by relevance vector regression. In *CVPR*, 2004.

[2] I. Akhter and M. J. Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *CVPR*, 2015.

[3] S. Amin, M. Andriluka, M. Rohrbach, and B. Schiele. Multi-view pictorial structures for 3d human pose estimation. In

BMVC, 2013.

[4] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014.

[5] C. S. Catalin Ionescu, Fuxin Li. Latent structured models for human pose estimation. In *ICCV*, 2011.

[6] D. A. Forsyth, O. Arikan, and L. Ikemoto. *Computational Studies of Human Motion: Tracking and Motion Synthesis*. Now Publishers Inc, 2006.

[7] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.

[8] M. Hofmann and D. M. Gavrila. Multi-view 3d human pose estimation in complex environment. *IJCV*, 2012.

[9] D. Hogg. Model-based vision: a program to see a walking person. *Image and Vision computing*, 1(1):5–20, 1983.

[10] P. Hu and D. Ramanan. Bottom-up and top-down reasoning with convolutional latent-variable models. *arXiv preprint arXiv:1507.05699*, 2015.

[11] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *PAMI*, 2014.

[12] H. Jiang. 3d human pose reconstruction using millions of exemplars. In *Pattern Recognition (ICPR), 2010 20th International Conference on*, 2010.

[13] S. Johnson and M. Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *Proceedings of the British Machine Vision Conference*, 2010.

[14] I. Kostrikov and J. Gall. Depth sweep regression forests for estimating 3d human pose from images. In *BMVC*, 2014.

[15] S. Li and A. B. Chan. 3d human pose estimation from monocular images with deep convolutional neural network. In *ACCV*, 2014.

[16] T. B. Moeslund and E. Granum. A survey of computer vision-based human motion capture. *Computer vision and image understanding*, 81(3):231–268, 2001.

[17] G. Mori and J. Malik. Recovering 3d human body configurations using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(7):1052–1062, 2006.

[18] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. *ECCV*, 2016.

[19] J. O'Rourke and N. I. Badler. Model-based image analysis of human motion using constraint propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (6):522–536, 1980.

[20] U. Rafi, J. Gall, and B. Leibe. A semantic occlusion model for human pose estimation from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2015.

[21] V. Ramakrishna, T. Kanade, and Y. Sheikh. Reconstructing 3d human pose from 2d image landmarks. In *ECCV*, 2012.

[22] G. Rogez and C. Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. *NIPS*, 2016.

[23] G. Shakhnarovich, P. Viola, and T. Darrell. Fast pose estimation with parameter-sensitive hashing. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 750–757. IEEE, 2003.

[24] J. Shotton, T. Sharp, A. Kipman, A. Fitzgibbon, M. Finocchio, A. Blake, M. Cook, and R. Moore. Real-time human pose recognition in parts from single depth images. *Communications of the ACM*, 2013.

[25] E. Simo-Serra, A. Ramisa, G. Alenyà, C. Torras, and F. Moreno-Noguer. Single image 3d human pose estimation from noisy observations. In *CVPR*, 2012.

[26] B. Tekin, A. Rozantsev, V. Lepetit, and P. Fua. Direct prediction of 3d body poses from motion compensated sequences. In *CVPR*, 2016.

[27] J. J. Tompson, A. Jain, Y. LeCun, and C. Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *Advances in neural information processing systems*, pages 1799–1807, 2014.

[28] C. Wang, Y. Wang, Z. Lin, A. L. Yuille, and W. Gao. Robust estimation of 3d human poses from a single image. In *CVPR*, 2014.

[29] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *CVPR*, 2016.

[30] X. K. Wei and J. Chai. Modeling 3d human poses from uncalibrated monocular images. In *ICCV*, 2009.

[31] H. Yasin, U. Iqbal, B. Kruger, A. Weber, and J. Gall. A dual-source approach for 3d pose estimation from a single image. In *CVPR*, 2016.

[32] H. Yub Jung, S. Lee, Y. Seok Heo, and I. Dong Yun. Random tree walk toward instantaneous 3d human pose estimation. In *CVPR*, 2015.

[33] X. Zhou, M. Zhu, S. Leonardos, K. G. Derpanis, and K. Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *CVPR*, 2016.